UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/751,090	12/29/2000	Richard W. Busser	4430-22	5019

22442 7590 10/27/2003

SHERIDAN ROSS PC
1560 BROADWAY
SUITE 1200
DENVER, CO 80202

EXAMINER	
MASKULINSKI, MICHAEL C	

ART UNIT	PAPER NUMBER
2184	6

DATE MAILED: 10/27/2003

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary	Application No.	Applicant(s)	
	09/751,090	BUSSER, RICHARD W.	
	Examiner	Art Unit	
	Michael C Maskulinski	2184	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If the period for reply specified above is less than thirty (30) days, a reply within the statutory minimum of thirty (30) days will be considered timely.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133).
- Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *29 December 2000* .

2a)☐ This action is **FINAL**. 2b)☒ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disp sition of Claims

4)☒ Claim(s) *1-19* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1-14 and 16-19* is/are rejected.

7)☒ Claim(s) *15* is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☒ The drawing(s) filed on *01 March 2001* is/are: a)☒ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

11)☐ The proposed drawing correction filed on _____ is: a)☐ approved b)☐ disapproved by the Examiner.

If approved, corrected drawings are required in reply to this Office action.

12)☐ The oath or declaration is objected to by the Examiner.

Pri rity under 35 U.S.C. §§ 119 and 120

13)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____ .

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

14)☐ Acknowledgment is made of a claim for domestic priority under 35 U.S.C. § 119(e) (to a provisional application).

a) ☐ The translation of the foreign language provisional application has been received.

15)☐ Acknowledgment is made of a claim for domestic priority under 35 U.S.C. §§ 120 and/or 121.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☐ Information Disclosure Statement(s) (PTO-1449) Paper No(s) _____ .
4) ☐ Interview Summary (PTO-413) Paper No(s). _____ .
5) ☐ Notice of Informal Patent Application (PTO-152)
6) ☐ Other:

Non-Final Office Action

Claim Objections

1. Claim 1 is objected to because of the following informalities: it is unclear as to what "substantially avoiding writing previously stored data" means. Writing data is either avoided or done. Based upon the Applicant's specification, the Examiner interprets the claim language to mean avoiding writing previously stored data means. Appropriate correction is required.

2. Claims 3, 5, 6, 7, 8, 10, 11 are objected to because of the following informalities: it is unclear as to what the limitation: "said using secondly step" refers to. The Examiner interprets the claim language as referring to the step in claim 1: "using secondly said array of storage devices..." even though this is the fourth step in the method of claim 1. Appropriate correction is required.

3. Claim 4 is objected to because of the following informalities: it is unclear as to what the limitation: "said making step" refers to. The Examiner interprets this claim language as referring to the step in claim 3: "making a determination related to being able to use said array of storage devices including said at least two thereof. Appropriate correction is required.

4. Claim 7 is objected to because of the following informalities: it is unclear as to what "said issuing step" includes. Does the "issuing step" include just issuing the trust array command or does it include issuing the trust array command and waiting for its completion. The Examiner has interpreted "said issuing step" to mean just issuing the trust array command. Appropriate correction is required.

5. Claim 13 is objected to because of the following informalities: it is unclear as to what "substantially avoiding restoring and/or reconstructing data and/or parity" means. Restoring and/or reconstructing data and/or parity is either avoided or done. Based upon the Applicant's specification, the Examiner interprets the claim language to mean avoiding restoring and/or reconstructing data and/or parity. Appropriate correction is required.

Claim Rejections - 35 USC § 102

6. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

> A person shall be entitled to a patent unless –
>
> (b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.
>
> (e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

7. Claims 1-8, 10-12, 14, 16-19 are rejected under 35 U.S.C. 102(e) as being anticipated by Renner, Jr., U.S. Patent 6,243,827 B1.

Referring to claim 1:

a. In column 6, lines 26-36, Renner, Jr. discloses a RAID 3 system comprising a single host computer, a RAID controller, and two tiers of 5 Direct Access Storage Device (DASD) units with two parity DASDs and two additional hot spare DASDs (using firstly an array of storage devices to conduct read/write

operations under control of at least a first controller including providing metadata to each of said storage devices of said array).

b. In column 12, lines 13-15, Renner, Jr. discloses a MEDIUM ERROR status reported by the storage unit (ascertaining that a failure has occurred).

c. In column 4, lines 30-47, Renner, Jr. discloses that catastrophic disk-array failures involve the failure of greater than one disk in a RAID 1-5 system, or any disk in a RAID 0 system (discontinuing use of at least two storage devices of said array related to conducting read/write operations based on the failure).

d. In column 11, lines 48-54, Renner, Jr. discloses that if the data was not successfully committed to the array even after retries and/or removal of a redundant disk from operation, the subroutine ADD BAD BLOCKS is called for the range of logical addresses described by the command (using secondly said array of storage devices after said discontinuing step while avoiding writing previously stored data and/or parity to said at least two storage devices that was present before the failure).

Referring to claim 2, in column 2, lines 3-11, Renner, Jr. discloses transient failures and controller failure. Further, in column 4, lines 53-56, Renner Jr. discloses a power failure (a transient failure related to a back plane).

Referring to claim 3, in column 4, lines 30-34, Renner, Jr. discloses that catastrophic disk-array failures involve the failure of greater than one disk in a RAID 1-5 system, or any disk in a RAID 0 system. In almost all cases, however, there remains the ability of the controller to communicate with at least one disk in the array (said using

secondly step includes making a determination related to being able to use said array of storage devices including said at least two thereof).

Referring to claim 4, in column 11, lines 31-36, Renner, Jr. discloses the process, which is performed when a write operation to the storage array has been completed, successfully or unsuccessfully, by the controller. In the case of a failure, this process will only be called after the number of retries prescribed by the errorhandling policies of the controller have been performed (said making step includes checking whether one or more of said storage devices is off-line).

Referring to claim 5, in column 4, lines 34-39, Renner, Jr. discloses that the present invention uses software and a small portion of each disk in the array to write a bad area table on each disk. The bad area table provides the logical address and length of the area in the array's logical space, which has been corrupted by physical damage on the media or other causes of write failure (said using secondly step includes modifying metadata on each of said storage devices).

Referring to claims 6 and 18, in column 9, lines 31-59, Renner, Jr. discloses that if BBM STAMP is valid, its timestamp is checked against the timestamp of CANONICAL STAMP (where the timestamp for a null stamp is defined to be older than any valid timestamp). If it is more recent than CANONICAL STAMP, CANONICAL STAMP is set to BBM STAMP. Control then passes to the next iteration of the loop. When there are no more disks to check, the CANONICAL STAMP is checked for a null value. If it is not null, then an appropriate stamp has been located, and all bad regions described in

CANONICAL STAMP are added to the BBM MAP (a partition age metadata field to be synchronized on each of said storage devices of said array).

Referring to claim 7, in column 9, lines 31-59, Renner, Jr. discloses that the presence of other disks is checked in a loop (said using secondly step includes determining whether each of said storage devices of said array is accessible after said issuing step).

Referring to claim 8, in column 11, lines 36-40, Renner, Jr. discloses that if the array has redundancy and errors occur on only one disk, the disk may be removed from operation prior to the invocation of this process, in which case the status of the write operation will be considered to be good (said using secondly step includes controlling re-use of said array when it is determined that no more than one of said storage devices of said array is off-line).

Referring to claim 10, in column 11, lines 36-40, Renner, Jr. discloses that if the array has redundancy and errors occur on only one disk, the disk may be removed from operation prior to the invocation of this process, in which case the status of the write operation will be considered to be good (said using secondly step includes controlling re-use of said array based on one of a user determination and an automatic determination independently of the user).

Referring to claim 11, in column 6, lines 59-67 continued in column 7, lines 1-6, Renner, Jr. discloses that the flow of data between host and disk array is indicated by the heavy line. Data is received from the host computer via the host SCSI bus into the SCSI input/output processor (SCSI IOP). The SCSI IOP initiates memory transactions

to or from the cache memory through the bridge chip, which bridges the system bus and the cache bus. A cache bus connects the bridge chip, cache memory, and the hardware control mechanism DMA Sync. The DMA Sync acts as a direct memory access (DMA) controller with the additional RAID-3 function of parity generation and checking and replacement of data with a hot spare. It also generates reads or writes to specific cache addresses and translates the data between the cache bus and the SCSI interface chip on the individual channel (said using secondly step includes generating a command by a host and transmitting said command to said first controller).

Referring to claim 12, it is inherent to have the read/write command, as discussed above, initiated manually by the user of the array.

Referring to claim 14:

a. In column 6, lines 26-36, Renner, Jr. discloses a RAID 3 system comprising a single host computer, a RAID controller, and two tiers of 5 Direct Access Storage Device (DASD) units with two parity DASDs and two additional hot spare DASDs (an array of storage devices relative to which read and write data transfers are conducted).

b. In column 6, lines 26-36, Renner, Jr. discloses a RAID 3 system comprising a single host computer, a RAID controller, and two tiers of 5 Direct Access Storage Device (DASD) units with two parity DASDs and two additional hot spare DASDs (a controller communicating with said array of storage devices for conducting read/write operations).

c. In column 6, lines 59-67 continued in column 7, lines 1-6, Renner, Jr. discloses that the flow of data between host and disk array is indicated by the heavy line. Data is received from the host computer via the host SCSI bus into the SCSI input/output processor (SCSI IOP). The SCSI IOP initiates memory transactions to or from the cache memory through the bridge chip, which bridges the system bus and the cache bus. A cache bus connects the bridge chip, cache memory, and the hardware control mechanism DMA Sync. The DMA Sync acts as a direct memory access (DMA) controller with the additional RAID-3 function of parity generation and checking and replacement of data with a hot spare. It also generates reads or writes to specific cache addresses and translates the data between the cache bus and the SCSI interface chip on the individual channel (a host communicating with said controller that makes requests related to data to be stored and data to be obtained from said array of storage devices).

d. In column 4, lines 30-47, Renner, Jr. discloses that catastrophic disk-array failures involve the failure of greater than one disk in a RAID 1-5 system, or any disk in a RAID 0 system. In almost all cases, however, there remains the ability of the controller to communicate with at least one disk in the array. The present invention uses software and a small portion of each disk in the array to write a bad area table on each disk. The bad area table provides the logical address and length of the area in the array's logical space, which has been corrupted by physical damage on the media or other causes of write failure. After a catastrophic failure of multiple disks, assuming at least one disk can be written

to, there will be a record of the failure on at least one disk (wherein said host is used in generating a trust array command related to updating metadata on each of said storage devices of said array after a fault occurs and after use of said array was discontinued due to the fault).

Referring to claim 16, in column 4, lines 30-47, Renner, Jr. discloses that catastrophic disk-array failures involve the failure of greater than one disk in a RAID 1-5 system, or any disk in a RAID 0 system. In almost all cases, however, there remains the ability of the controller to communicate with at least one disk in the array. The present invention uses software and a small portion of each disk in the array to write a bad area table on each disk. The bad area table provides the logical address and length of the area in the array's logical space, which has been corrupted by physical damage on the media or other causes of write failure. After a catastrophic failure of multiple disks, assuming at least one disk can be written to, there will be a record of the failure on at least one disk. The task of repairing the array is greatly simplified because all bad regions of the array can be easily identified (said trust array command is generated independently of any reconstruction and/or restoration of said array).

Referring to claim 17, in column 11, lines 31-40, Renner, Jr. discloses that in the case of a failure, if the array has redundancy and errors occur on only one disk, the disk may be removed from operation prior to the invocation of this process (said host controls discontinuing use of said array of storage devices based on the fault). Further, in column 11, lines 48-54, Renner, Jr. discloses that if the data was not successfully committed to the array even after retries and/or the removal of a redundant disk from

operation, the subroutine ADD BAD BLOCKS is called for the range of logical addresses described by the command (and subsequently said host receives an input that is used in generating said trust array command).

Referring to claim 19, in column 11, lines 31-36, Renner, Jr. discloses that in the case of a failure, the process will only be called after the number of retries prescribed by the error handling policies of the controller have been performed (a determination is made by at least one of said controller and said host related to whether one or more of said storage devices is off-line before said trust array command is generated). Further, in column 4, lines 30-47, Renner, Jr. discloses that after a catastrophic failure of multiple disks and the bad area table is written, the system determines which disks are accessible (a determination is made by at least one of said host and said controller related to whether each of said storage devices of said array is accessible after said trust array command is generated).

8. Claims 1 and 13 are rejected under 35 U.S.C. 102(b) as being anticipated by Jones, U.S. Patent 5,479,653.

Referring to claim 1:

a. In column 4, lines 54-55, Jones discloses a disk array controller and disk drives (using firstly an array of storage devices to conduct read/write operations under control of at least a first controller). Further, in column 6, lines 48-51, Jones discloses that if a RAID 5 implementation is used instead of a RAID 4, then the virtual disk driver also performs parity distribution operations to distribute

parity among each of the three drives in the three drive RAID 5 set (providing metadata to each of said storage devices of said array).

b. In column 6, lines 24-26, Jones teaches ascertaining that a failure has occurred.

c. In column 3, lines 13-20, Jones discloses that the disk array system of the present invention begins operation with an array configuration that implements a maximum of data redundancy using all of the available disk drives. As drive failures occur, the disk array system automatically reconfigures one or more of the drives to other RAID configurations which utilize less data redundancy and thus require a lesser number of drives (discontinuing use of at least two of said storage devices of said array related to conducting read/write operations based on the failure).

d. In column 3, lines 16-20, Jones discloses that as failures occur, the disk array system automatically reconfigures one or more of the drives to other RAID configurations which utilize less data redundancy and thus require a lesser number of drives (using secondly said array of storage devices after said discontinuing step while substantially avoiding writing previously stored data and/pr parity to said at least two storage devices that was present before the failure).

Referring to claim 13, in column 3, lines 16-20, Jones discloses that as failures occur, the disk array system automatically reconfigures one or more of the drives to other RAID configurations which utilize less data redundancy and thus require a lesser

number of drives (said substantially avoiding writing previously stored data and/or parity includes substantially avoiding restoring and/or reconstructing data and/or parity).

Claim Rejections - 35 USC § 103

9. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

10. Claim 9 is rejected under 35 U.S.C. 103(a) as being unpatentable over Renner, Jr., U.S. Patent 6,243,827 B1 as applied to claim 1 above, and further in view of Stephenson, U.S. Patent 6,353,895 B1.

Referring to claim 9, in the Abstract, Renner, Jr. discloses a software-based method for facilitating the recovery of a RAID storage system from the simultaneous failure of two or more disks. However, Renner, Jr. doesn't explicitly disclose allowing data and/or parity to be read by said first controller when more than one of said storage devices is off-line and reading said data and/or parity from said storage devices of said array that are on-line. In column 2, lines 37-49, Stephenson discloses that one embodiment includes a redundant array of independent disk drives that provides one-drive and two-drive fault tolerance. Data recovery from a one or two drive failure is accomplished by using a two-dimensional XOR parity arrangement. It would have been obvious to one of ordinary skill at the time of the invention to include the two-drive fault tolerance and parity arrangement of Stephenson into the system of Renner, Jr. A

person of ordinary skill in the art would have been motivated to make the modification because *this parity arrangement uses less storage than mirroring when the number of total drives is greater than four* (see Stephenson: Abstract).

Allowable Subject Matter

11. Claim 15 is objected to as being dependent upon a rejected base claim, but would be allowable if rewritten in independent form including all of the limitations of the base claim and any intervening claims.

Conclusion

12. The prior art made of record and not relied upon is considered pertinent to applicant's disclosure.

U.S. 2002/0016942 A1	MacLaren et al.
U.S. 6,557,123 B1	Wiencko, Jr. et al.
U.S. 6,530,035 B1	Bridge
U.S. 6,408,406 B1	Parris

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Michael C Maskulinski whose telephone number is (703) 308-6674. The examiner can normally be reached on Monday-Friday 9:30-6:00.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Robert W Beausoliel can be reached on (703) 305-9713. The fax phone number for the organization where this application or proceeding is assigned is (703) 872-9306.

Any inquiry of a general nature or relating to the status of this application or proceeding should be directed to the receptionist whose telephone number is (703) 305-3900.

MM

ROBERT BEAUSOLIEL
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100